Heterogeneous Innovation and the Antifragile Economy*

Gustavo Manso [a], Benjamin Balsmeier [b], and Lee Fleming [a]

a) University of California, Berkeley, USA
b) University of Luxembourg

April 2019

Abstract: Schumpeter (1939) claims that recessions are periods of "creative destruction", concentrating innovation that is useful for the long-term growth of the economy. However, previous research finds that standard measures of firms' innovation, such as R&D expenditures or patenting, concentrate in booms. We argue that these standard measures do not capture shifts in firms' innovative search strategies. We introduce a model of firms' choice between exploration vs. exploitation over the business cycle and find evidence that firms shift towards exploration during contractions and exploitation during expansions. Results are stronger for firms in more cyclical industries and with weaker financial constraints.

Keywords: Exploration, Exploitation, Patents, Innovation, Business Cycles, Macroeconomic Risk, Productivity, Growth

JEL Codes: O31, O32

1. Introduction

Antifragility describes systems that improve in capability when exposed to volatility and negative shocks.[1] An antifragile economy is thus one that becomes stronger when exposed to macroeconomic fluctuations. Such an idea is at odds with macroeconomic policy whose goal is stability. In fact, it suggests that the recent great moderation period from the mid-80s to 2007, which witnessed decreased macroeconomic volatility, may have actually been ultimately detrimental to the economy, exposing it to bigger risks that culminated in the collapse of the financial system and stagnation of productivity gains.

Schumpeter (1939) supports the view that long-term productivity may benefit from macroeconomic fluctuations. According to him, recessions are times of creative destruction, in which increased innovation fuels enhancements in productivity and the retirement of old technologies. A large body of theoretical work – including Cooper and Haltingwanger (1993), Caballero and Hammour (1994), Aghion and Saint-Paul (1998), and Canton and Uhlig (1999) – has formalized Schumpeter's thesis. This literature typically builds upon the simple idea that the opportunity cost of firms' innovative activities, i.e. the foregone sales that could have been achieved instead, drops in recessions. Stated another way, during recessions, firms should focus on long-run investments since expected profits in the short run are low anyways.

A number of famous anecdotes about firms' innovations can be adduced to support these arguments. Dupont's dominance in the mid 20th century can be directly traced to the inventions from Wallace Caruthers's lab and others during the depression, including neoprene (1930), nylon (1935), teflon (1938), and polyester (1941). Karl Jansky at Bell Labs discovered radio waves in 1931, in the process of tracking down sources of radio static. Following WWII and the accompanying downturn, Percy Spencer invented the microwave oven in 1946, and in 1947 Shockley, Bardeen, and Brattain at Bell Labs invented the transistor, which in turn enabled the electronics, information, and artificial intelligence revolutions.

Despite the plausible models and salient anecdotes, much systematic evidence suggests that firms do not take the opportunity to replenish the stock of productivity enhancing innovations

---

[1] See, for instance, Taleb (2012) for a discussion of the concept. A classic example of antifragility is how physical exercise, which creates oxidative stress and distresses muscle fibers, followed by periods of rest enhances strength and overall health.

during downturns. Typically measured by R&D expenditures and patents, most empirical work to date finds a procyclical bias for innovative activities (Griliches 1990, Geroski and Walters 1995, Fatas 2000, Rafferty 2003, Walde and Woitek 2004, and Comin and Gertler 2006, Kopytov, Roussanov, and Taschereau-Dumouchel, 2018). Field (2003) offers rare evidence in favor of the Schumpeterian hypothesis with time series measures of productivity. Yet most of the empirical work presents a conundrum; based on measures of R&D spending and patent counts, the data clearly reject the theoretical predictions of countercyclical innovation.

A variety of explanations have been proposed to explain the contrary evidence, for example, that firms invent in downturns but delay the commercialization of their inventions until demand increases (Schleifer 1986, Francois and Lloyd-Ellis 2003), that fear of appropriation encourages pro-cyclical innovation (Barlevy 2007), that credit constrained firms are less likely to invest in counter-cyclical innovation (Aghion et al. 2012), that pro-cyclical innovation is more likely in industries with faster obsolescence and weak intellectual property protection (Fabrizio and Tsolmon 2014), and that inventors become less productive during downturns, due to a deterioration in their household balance sheet (Bernstein, McQuade, and Townsend 2018).

To resolve this conundrum, we model innovative search as a tension within firms between exploration (the pursuit of novel to the firm approaches) versus exploitation (the refinement of existing technology that is known to the firm). We observe this tension empirically with a patent-based measure of technological proximity across time within each firm.

The model begins with the assumption that innovation results from experimentation with new ideas (Arrow 1969). The central tension that arises in experimentation lies between exploration and exploitation (March 1991). Exploration involves search, risk-taking and experimentation with new technologies or new areas of knowledge. Exploitation, on the other hand, is the refinement of existing and familiar technologies. Exploration is more expensive due to an increased probability of failure and the learning that it requires to commercialize new technologies. Because the opportunity cost of exploratory activities – the additional output or sales that could have been achieved instead by a slightly refined product – is lower in recessions, firms have incentives to undertake such activities in downturns. At the same time, during booms, firms have incentives to engage in exploitation, to avoid losing profits from the high sales of its current products. As a consequence, the model predicts that exploration is countercyclical while exploitation is procyclical.

The model and results are related to the literature on incentives for innovation (e.g. Holmstrom 1989; Aghion and Tirole 1994. Modelling the innovation process as a simple bandit problem, Manso (2011) finds that tolerance for early failure and reward for long-term success is optimal to motivate exploration. A similar principle operates in our model. During recessions, profit is low regardless of the action pursued, and thus the firm tolerates early failures. Moreover, future profits look more promising than the present, and thus there will be rewards for long-term success. In contrast to models which focus on innovation outcomes (Akcigit and Kerr, 2016), we focus on heterogeneity in the search process. Our model starts from the perspective of an individual firm and asks when it is more or less likely to leave already known to the firm paths.

To measure exploration and exploitation we still rely on patent data, however, we differentiate between patents filed in new to the firm technology classes and patents filed in known to the firm technology classes. We observe the distribution of the number of patents (in year of application) per technology class and firm. Building on Jaffe (1989) and Bloom et al. (2013), we then calculate the similarity between the distribution of patents across technology classes applied by a given firm in year t and the same firm's prior distribution of patents across technology classes. The technological profiles of firms that exploit will look more similar to their past profiles; those that explore will look different from year to year. Consistent with the model prediction, profiles become more similar during expansions and less similar during contractions.

Our main contribution is to break down firms' innovation and search strategies into exploration and exploitation over the business cycle and provide supporting evidence of shifts in firms' innovative search over the cycle. Data come from the joint availability of Compustat and patent observations for publicly traded firms from 1958 through 2008. Using this more nuanced view of innovation and within firm search strategy, we predict and find that innovative exploration is countercyclical while exploitation is procyclical. Moreover, we predict and find stronger results for more cyclical and less financially constrained firms. The results are robust to a variety of estimations, alternative measures, and data cuts.

These results suggest that changes within firms' search strategies can bolster economic antifragility and perhaps cast a more positive view of the welfare effects of macroeconomic fluctuations. If negative economic shocks indeed encourage growth-enhancing exploration, economic recessions would tend to be shorter and less persistent than they would be otherwise.

4

Cyclical fluctuations could contribute positively to welfare if they helped balance exploration and exploitation within firms. This positive contribution might be even more important, if there exists an inherent bias towards exploitation, for example, due to the imperfect protection of property rights, or the difficulty of commercializing new technologies and appropriating their profits for the inventing firm. If the normal balance tilts away from the optimal balance of exploration and exploitation, then macroeconomic fluctuations might perform an important function in renewing each firm's stock of innovations that ultimately fuel productivity improvements and economic growth.

2. Models and Predictions

2.1. The Base Model

We introduce a model of exploration and exploitation over the industry business cycle. The model is based on the simple two-armed bandit problem studied in Manso (2011), but incorporates macroeconomic shocks.

The economy exists for two periods. In each period, the representative firm in the economy takes either a well-known or a novel action. The well-known action has a known probability p of success (S) and $1 - p$ of failure (F) with $S > F$. The novel action has an unknown probability q of success and $1 - q$ of failure (F). The only way to learn about q is by taking the novel action. The expected probability of success when taking the novel action is $E[q]$ when the action is taken for the first time, $E[q|S]$ after experiencing a success with the novel action, and , $E[q|F]$ after experiencing a failure with the novel action. From Bayes' rule, $E[q|F] < E[q] < E[q|S]$.

We assume that the novel action is of exploratory nature. This means that when the firm experiments with the novel action, it is initially not as likely to succeed as when it conforms to the conventional action. However, if the firm observes a success with the novel action, then the firm updates its beliefs about the probability q of success with the novel action, so that the novel action becomes perceived as better than the conventional action. This is captured as follows:

$$E[q] < p < E[q|S].$$

The macroeconomic state m can be either high (H) or low (L). If the macroeconomic state is currently m it remains in the same state next period with probability μ. Alternatively, it

5

transitions into the other state n next period. Industry demand in macroeconomic state m is d_m with $d_H > d_L$. Given the macroeconomic state m, firm profit in each period is given by $d_m S$ in case of success and $d_m F$ in case of failure.

For simplicity, we assume risk-neutrality and a discount factor of δ. There are only two action plans that need to be considered. The first relevant action plan, exploitation, is to take the well-known action in both periods. This action plan gives the payoff $\pi(m, exploit)$ if the macroeconomic state is m:

$$p d_m S + (1 - p) d_m F + \delta \mu (p d_m S + (1 - p) d_m F) + \delta(1 - \mu)(p d_n S + (1 - p) d_n F)$$

The other relevant action plan, exploration, is to take the novel action in the first period and stick to it only if success is obtained. This action plan gives the payoff $\pi(m, explore)$ if the macroeconomic state is m:

$$\begin{aligned} E[q] d_m S &+ (1 - E[q]) d_m F \\ &+ \delta\mu \left(E[q](E[q|S] d_m S + (1 - E[q|S])) d_m F \right) \\ &+ (1 - E[q])(p\, d_m S + (1 - p) d_m F) \right) + \delta(1 \\ &- \mu)(E[q](E[q|S] d_n S + (1 - E[q|S]) d_n F) + (1 \\ &- E[q])(p d_n S + (1 - p) d_n F)) \end{aligned}$$

The total payoff from exploration is higher than the total payoff from exploitation if:

$$E[q] \geq \frac{d_m p}{d_m(1 + \delta\,(E[q|S] - p)\mu) + d_n\,\delta\,(E[q|S] - p)(1 - \mu)}\; p \qquad (1)$$

If the firm tries the novel action, it obtains information about q. This information is useful for the firm's decision in the second period, since the firm can switch to the conventional action if it learns that the novel action is not worth pursuing. The fraction multiplying p in the inequality above is less than 1. Therefore, the firm may be willing to try the novel action even though the initial expected probability $E[q]$ of success with the novel action is lower than the probability p of success with the conventional work method.

Proposition 1: Firms are more prone to explore in recessions than in booms.

Proof: The coefficient multiplying p on the right-hand side of equation (1) is increasing in d_m and decreasing in d_n. Since $d_H > d_L$, the firm is more prone to explore in bad times ($m = L, n = H$) than in a good times ($m = H, n = L$). ∎

The intuition for the result is that in a recession, the future is more important than the present, since current industry demand is low. Therefore, the firm is more forward-looking and is willing to explore for a larger set of opportunities.

2.2.Industry Cyclicality

How do results vary with industry cyclicality? More cyclical industries are those that respond more to the macroeconomic state (higher d_m and lower d_n). The following proposition studies this comparative statics.

Proposition 2: The innovation strategies of firms in cyclical industries are more sensitive to business cycles.

Proof: Since the coefficient multiplying p on the right-hand side of equation (1) is increasing in d_m, decreasing in d_n, and $d_H > d_L$, more cyclical firms are more prone to exploration than less cyclical firms during recessions. Conversely, more cyclical firms are less prone to exploration than less cyclical firms during booms. ∎

The intuition is that, for more cyclical firms, fluctuations caused by the business cycle are exaggerated. This amplifies the dependence of innovation strategy on the business cycle, derived in Proposition 1.

2.3. Financial Constraints

We extend the model to allow for financial constraints. To capture financial constraints we allow the discount rate to differ depending on the state of the economy. Because financial constraints are more likely to bind during recessions, we assume that for financially constrained firms the discount factor δ_L during bad times is lower than the discount factor $\delta_H = \delta$ during good times.

Again, there are only two action plans that need to be considered. The first relevant action plan, exploitation, is to take the well-known action in both periods. This action plan gives the following payoff $\pi_m(exploit)$ if the macroeconomic state is m:

$$p\mathrm{d_m}S + (1-p)d_m F + \delta_m \mu \, (pd_m S + (1-p)d_m F) + \delta_m(1-\mu)(pd_n S + (1-p)d_n F)$$

The other relevant action plan, exploration, is to take the novel action in the first period and stick to it only if success is obtained. This action plan gives the following payoff $\pi_m(explore)$ if the macroeconomic state is m:

$$E[q]d_m S + (1-E[q])d_m F + \delta_m \mu \left(E[q](E[q|S]d_m S + (1-E[q|S]))d_m F) + (1-E[q])(p\,d_m S + (1-p)d_m F) \right) + \delta_m(1-\mu)(E[q](E[q|S]d_n S + (1-E[q|S])d_n F) + (1-E[q])(pd_n S + (1-p)d_n F))$$

The total payoff from exploration is higher than the total payoff from exploration if:

$$E[q] \geq \frac{d_m p}{d_m(1 + \delta_m \, (E[q|S] - p)\mu)) + d_n \, \delta_m \, (E[q|S] - p)(1-\mu)} \, p$$

As before, the fraction multiplying p in the inequality above is less than 1. Therefore, the firm may be willing to try the novel action even though the initial expected probability $E[q]$ of success with the novel action is lower than the probability p of success with the conventional action.

Proposition 3: The innovation strategies of financially constrained firms are less sensitive to business cycles.

Proof: The coefficient multiplying p on the right-hand side of the inequality above is decreasing in δ_m. Because $\delta_L < \delta_H = \delta$, the innovation strategy of a financially constrained firm is less sensitive to business cycles. ∎

The intuition is that financially constrained firms discount the future more during recessions, offsetting the positive impact of macroeconomic shocks on exploration.

2.4. Antifragility

While we cannot test the prediction empirically, we extend the model and consider how economic welfare might respond to an increase in macroeconomic volatility. For that, we consider mean preserving spreads in $\{d_H, d_L\}$. The next proposition studies the effects of economic fluctuations on economic welfare.

Proposition 4: Welfare is higher in an economy with mean preserving macroeconomic fluctuations than in a stable economy.

Proof: The result follows from Jensen's inequality. In an economy with fluctuations, the representative firm can achieve at least the same profit as in a stable economy by following the optimal stable economy strategy regardless of the macroeconomic state:

$$\frac{1}{2}\pi(H, exploit) + \frac{1}{2}\pi(L, exploit) = (1 + \delta)(pH + (1 - p)L)$$

$$\frac{1}{2}\pi(H, explore) + \frac{1}{2}\pi(L, explore)$$
$$= E[q]S + (1 - E[q])F$$
$$+ \delta\left(E[q](E[q|S]\,S + (1 - E[q|S]))F) + (1 - E[q])(p\,S + (1 - p)F)\right)$$

Strict inequality holds if the optimal strategy (exploration vs exploitation) with fluctuations depends on the macroeconomic state. ∎

The economy is thus antifragile in the sense that it benefits from macroeconomic volatility. With macroeconomic fluctuations the firm can tailor its innovation strategy to the macroeconomic state, exploring during recessions and exploiting during booms. This flexibility leads to more creative destruction and higher welfare.

Another way to grasp the intuition behind the result is to note that the investment technology in this economy is a real option. The firm can adjust its strategy to the realization of the state of the economy. Since volatility typically increases option value, the economy benefits from macroeconomic fluctuations.

The economy in our base model is antifragile and benefits from macroeconomic fluctuations. Suppressing those fluctuations may reduce welfare. Obviously, if we allowed for risk-aversion some degree of (but not absolute) macroeconomic stability would be desirable.

3. Empirical Methodology

In order to empirically distinguish firms in any given year based on their relative focus on exploitation of known to the firm technologies, versus exploration of new to the firm technologies (otherwise referred to as a firm's *innovation search*), we draw on the original technology classes that USPTO examiners assigned to each patent.[2] Our measure examines the degree of overlap between patents granted to the firm in year t and the existing patent portfolio held by the same firm up to year $t - 1$. In particular, we employ the following variant of Jaffe's (1989) technological proximity measure to estimate similarity in technological space of firm i's patents applied in year t and its pre-existing patents applied for between $t - 5$ and $t - 1$, using patent counts per USPTO three-digit technology classes k:

$$innovative\ search_{i,t} = 1 - \frac{\sum_{k=1}^{K} f_{i,k,t} f_{i,k,t-1}}{\left(\sum_{k=1}^{K} f_{i,k,t}^2\right)^{\frac{1}{2}} \left(\sum_{k=1}^{K} f_{i,k,t-1}^2\right)^{\frac{1}{2}}} \tag{1}$$

where $f_{i,k,t}$ is the fraction of patents granted to firm i in year t that are in technology class k such that the vector $f_{i,t} = (f_{i,1,t} \dots f_{i,K,t})$ locates the firm's year t patenting activity in K-dimensional technology space.[3] *Innovative Search$_{i,t}$* is basically one minus the cosine angle between both vectors and would be one for a given firm-year when there is no overlap of patents' technology classes in year t compared to the previous five years; *Innovative Search$_{i,t}$* will equal zero when the distribution of firm i's patents applied in a given year is identical to patents accumulated in previous five years. When firms search for new technologies extensively, i.e. patent *only* in new to the firm technology classes, the measure

---

[2] If there is more than one technology class assigned to a patent we take the first one mentioned on the patent grant.
[3] Results are robust to taking all prior patents applied by the given firm into account, changing the threshold value from 5 to 10 years, and applying a 15% depreciation rate to a firm's past patent stock per technology class when calculating the innovative search measure.

would be one. Therefore, we classify firms as being relatively more focused on exploration/(exploitation) when they have a high/(low) *Innovative Search*$_{i,t}$ score. Bloom et al. (2013) use a similar approach to measure technological similarity across firms rather than within firms over time. They also study and discuss alternative measures of technological similarity in detail but find little differences in their results.

We follow Fabrizio and Tsolmon (2014) in adapting the classic patent production model (Hall, Griliches, & Hausman, 1986, and Pakes & Griliches, 1980) to estimate the effect of changes in industry demand on within firm changes in innovative search. Specifically, we estimate the following equation in OLS[4]:

$$IS_{it} = \alpha_0 + \beta_1 D_{kt-1} + \beta_2 X_{it-1} + f_i + \delta_t + \varepsilon_{it}, \qquad (2)$$

where IS_{it} is the innovative search focus of firm *i* in year *t*, D_{kt-1} is the output in industry *k* in year *t-1*, X_{it-1} is a vector of one-year lagged firm level controls, and f_i controls for time-invariant unobserved firm characteristics. Besides reducing endogeneity concerns, the latter resembles the theoretical prediction of shifts towards more or less exploration (exploitation) *within* firms. δ_t denotes a full set of year fixed effects that absorb aggregate changes in industry demand due to varying macroeconomic conditions, and ε_{it} is the error term.

If industry specific output strongly co-varies with the macro economy, however, this may leave little unique variation to identify how firms change their innovative search in response to changes in macroeconomic conditions. We thus follow Barlevy (2007) and estimate a model without time fixed effects in addition.[5] This empirical model should reflect firms' reactions to macroeconomic shocks more accurately, however, it has the unavoidable downside of being potentially confounded by aggregate changes in patent policies or subsidies that affect all firms and industries at a given point of time.

As in Fabrizio and Tsolmon (2014) the vector X_{it-1} contains controls for R&D spending, sales, employment and property, and plant and equipment per firm. Controlling for firms' sales should reduce concerns that the output measure captures the firm specific change in sales, and controlling for employment should capture firm size variation over the business cycle, and

---

[4] Alternatively estimating a quasi-fixed effects Tobit model in the spirit of Chamberlain (1986) and proposed by Wooldridge (2002, p. 538f.) reveals qualitatively the same results.
[5] Alternatively, we also estimated models where δ_t is replaced by linear or log-linear cycle trend, drawing on the NBER US Business Cycle Expansions and Contractions data, where the trend variable take value zero in recession periods and values 1, 2, …, N, for the first, second, …, and Nth year of each expansion period. Results remain unchanged. The trend itself is significantly positive, and taking just recession dummies instead of a trend indicate increases in exploration during recession periods.

property, plant and equipment should capture changes in physical capital. A positive (negative) estimated coefficient on D_{kt-1} would indicate that, controlling for any change in R&D spending, firms focus more on exploration (exploitation) when industry output increases. Observed changes in innovative search are thus not just driven by the procyclical changes in R&D as shown in Barlevy (2007). The lagged aggregated industry output, net of the focal firms' sales, should also be invulnerable to reverse causality and largely exogenous to firm specific choice, thus reducing concerns about potential endogeneity biases. Please see the Appendix for discussion for a variety of models that investigate robustness and endogeneity concerns.

4. Data

The empirical analysis is based on the joint availability of firm level data from three sources: 1) public US based firms in Compustat, 2) disambiguated patent assignee data from Kogan et al. (2017), the United States Patent and Trademark Office, and the Fung Institute at UC Berkeley (Balsmeier et al. 2018), and 3) the NBER-CES Manufacturing Industry Database (Bartelsman & Gray, 1996). We build firm level patent portfolios by aggregating eventually granted US patents from 1958 (first year of availability of the NBER-CES industry data) through 2008 inclusive. Kogan et al. (2017) provide data on patents granted through 2010, however, we truncate the sample at 2008 because patent pendency averages three years, and we model patents at their time of application, not grant. As we base our analysis on measures that have no obvious value in case of non-patenting activity or first time patenting activity, we only include firms in the analysis that applied for at least one patent in a given year, and patented at least once in any previous year, taking all patents granted to a given firm back to 1926 into account. The match with the NBER-CES database reduces the sample to manufacturing industries. Firms in manufacturing account for about 70 to 80% of the economy wide R&D spending since 1990 and about 90% beforehand (Barlevy, 2007). Finally, we restrict the sample to firms that we observe at least twice and have non-missing values in any control variable. The final dataset is an unbalanced panel of 21,051 firm year observations on 1,893 firms in 124 manufacturing industries, observed between 1958 and 2008.

Following Barlevy (2007), we measure industry output at the 4-digit SIC industry level.[6] We take the same measure of industry output as our predecessors, namely the value added and material costs per industry, deflated by each industries' shipments deflator as provided by the

---

[6] Results are robust to higher aggregation to the 3-digit SIC industry level. This level is less precise but also less likely to pick any unobserved time-varying change in firm characteristics.

NBER-CES database. R&D expenses, sales and capital are deflated by the official IMF US price inflation index. Table 1 presents summary statistics. Please see the Appendix for a series of measure and modeling robustness estimations.

Table 1 – Summary statistics

Variable	N	mean	Median	sd	min	max
Innovative Search	21051	0.40	0.32	0.32	0.00	1.00
Log(R&D)$_{t-1}$	21051	2.03	1.91	1.93	-4.90	8.80
Log(Sales)$_{t-1}$	21051	12.38	12.49	2.33	0.81	18.97
Log(Employees)$_{t-1}$	21051	1.69	1.41	1.36	0.00	6.78
Log(Capital)$_{t-1}$	21051	4.05	3.97	2.37	-5.18	10.97
Log(Output)$_{t-1}$	21051	9.44	9.28	1.61	2.97	15.01

Notes: This table reports summary statistics of variables used in the study. Sample covers all public US firms covered by Compustat that patented at least twice between 1958 and 2008. Innovative search is the technological proximity between the patents filed in year *t* to the existing patent portfolio held by the same firm from year *t*-5 up to year *t*-1, calculated according to Jaffe (1989). R&D, sales and capital (property, plant, and equipment) are from Compustat and deflated by the IMF price index. Output is value added and material costs per SIC 4-digit manufacturing industry, deflated by each industries' shipments deflator as provided by the NBER-CES database.

Table 2 - Correlation matrix

		(1)	(2)	(3)	(4)	(5)	(6)
(1)	Innovative Search	1.000					
(2)	Log(R&D)$_{t-1}$	-0.318	1.000				
(3)	Log(Sales)$_{t-1}$	-0.095	0.479	1.000			
(4)	Log(Employees)$_{t-1}$	-0.129	0.489	0.899	1.000		
(5)	Log(Capital)$_{t-1}$	-0.126	0.520	0.928	0.906	1.000	
(6)	Log(Output)$_{t-1}$	-0.119	0.272	0.146	0.097	0.162	1.000

Notes: This table reports pairwise correlations of the log-transformed variables used in the study. Sample covers all public US firms covered by Compustat that patented at least twice between 1958 and 2008. Innovative search is the technological proximity between the patents filed in year *t* to the existing patent portfolio held by the same firm from year *t*-5 up to year *t*-1, calculated according to Jaffe (1989). R&D, sales and capital (property, plant, and equipment) are from Compustat and deflated by the IMF price index. Output is value added and material costs per SIC 4-digit manufacturing industry, deflated by each industries' shipments deflator as provided by the NBER-CES database.

4.1 Baseline results

We first confirm the pro-cyclicality of R&D spending (Barlevy, 2007), and patenting (Fabrizio and Tsolmon, 2014), with our longer time series (though smaller dataset, due to the patenting criterion for inclusion). As can be seen in Table 3, columns (a) and (b) for R&D spending, and (c) and (d) for patenting, these measures correlate positively with increases in aggregate output

per industry. As expected, and similar to the prior results, the impact weakens if we control for changes in the macro economic conditions that affect all firms and industries in the same way through the inclusion of year fixed effects. Table 3, columns (e) and (f), show the results of estimating our main model as introduced above, first without (e) and then with time fixed effects (f). The negative coefficients for the output variable supports the prediction of our theoretical model - that firms tend to explore less, i.e. search amongst known technologies, the better the economic conditions.

The magnitude of the effects are not only statistically but also economically significant. A one standard deviation increase in output corresponds to a 0.31 (model a) (0.10 [model b]) standard deviation increase in R&D spending, a 0.15 (model c) (0.26 [model d]) standard deviation increase in patenting, and a -0.18 (model e) (-0.12 [model f]) standard deviation decrease in innovative search/exploration.

Table 3 – Industry growth, R&D, patents and innovative search

	R&D spending		Patents		Innovative search	
	a	b	c	d	e	f
Log(R&D)$_{t-1}$			-0.035*	0.045**	-0.002	-0.004
			(0.019)	(0.022)	(0.003)	(0.003)
Log(Sales)$_{t-1}$	0.229***	0.153***	-0.080***	-0.009	0.009*	0.011*
	(0.071)	(0.039)	(0.024)	(0.021)	(0.005)	(0.006)
Log(Employees)$_{t-1}$	0.361**	0.303***	0.469***	0.478***	-0.040***	-0.056***
	(0.161)	(0.090)	(0.076)	(0.073)	(0.014)	(0.014)
Log(Capital)$_{t-1}$	0.430***	0.272***	0.113***	0.104***	-0.028***	-0.018*
	(0.054)	(0.039)	(0.043)	(0.037)	(0.010)	(0.009)
Log(Output)$_{t-1}$	0.369***	0.117***	0.119*	0.223***	-0.036***	-0.026***
	(0.119)	(0.026)	(0.064)	(0.046)	(0.008)	(0.006)
N	21051	21051	21051	21051	21051	21051
Year fixed effects	No	Yes	No	Yes	No	Yes
Firm fixed effects	Yes	Yes	Yes	Yes	Yes	Yes
R^2	0.813	0.857	0.754	0.779	0.466	0.474

Notes: This table presents OLS regression of firms' log(R&D spending), a and b, log(no. patents +1), c and d, and innovative search, e and f, defined as the technological proximity between the patents filed in year *t* to the existing patent portfolio held by the same firm from year *t*-5 up to year *t*-1, calculated according to Jaffe (1989). Standard errors clustered at the industry level appear in parentheses. ***, ** and * indicate a significance level of 1%, 5%, and 10%, respectively.

4.2 Pro-cyclical industries

Our theory further implies that the decreasing focus on exploration over the business cycle is stronger for firms that are active in particularly pro-cyclical industries as opposed to less

cyclical industries. To test this prediction empirically we build on Barlevy (2007) by measuring each industries' cyclicality with the correlation of publicly traded firms' stock market value with the industries' overall growth as measured by the NBER-CES. The idea is that the stock price reflects the discounted value of future dividends of publically traded firms. Specifically, we took all domestic firms in each industry at the 2-digit SIC level and regressed the growth rate in real stock prices per firms in a given industry on the real industry growth and a constant.[7] The coefficients on real growth from these regressions, named $\hat{\beta}_{stock}$, then reflect the degree to which stock market values per industry co-vary with the business cycle. Barlevy (2007) ran qualitatively the same regressions but exchanged the firms' market value growth with R&D growth, to derive a corresponding measure of how much R&D investments co-vary with the business cycle per industry. We calculate the same but use the growth in firms' innovative search score instead of R&D growth to derive our measure $\hat{\beta}_{isearch}$ of the pro-cyclicality of each industries' innovative search focus.

With these measures we regressed $\hat{\beta}_{isearch}$ on $\hat{\beta}_{stock}$, yielding: $\hat{\beta}_{isearch} = -0.048 - \underset{SE=0.0134}{\underbrace{0.765}} \times \hat{\beta}_{stock}$.[8] This equation is consistent with our theory predicting stronger decreases in innovative search (exploration) over the business cycle, the more pro-cyclical the industry.

We also test this prediction by estimating a slightly abbreviated version of our baseline model:

$$IS_{it} = \alpha_0 + \beta_1 D_{kt-1} + \beta_2 X_{it-1} + \beta_3 D_{kt-1} \times Cyc_k + f_i + \delta_t + \varepsilon_{it}, \qquad (3)$$

where we keep everything as introduced above but add an interaction of industry demand D_{kt} and a dummy for strong industry cyclicality Cyc_k, i.e. a $\hat{\beta}_{stock}$ value above the median. For easier comparison we keep $D_{kt-1} \times Cyc_k$ where Cyc_k is equal to one and replace all values of D_{kt-1} with zero if Cyc_k is equal to zero such that the size of β_1 is the estimated elasticity of demand and innovative search in weakly pro-cyclical or counter cyclical industries and β_3 is the estimated elasticity of demand and innovative search in strongly pro-cyclical industries. Note that the main effect of Cyc_k is fully absorbed by f_i. A larger estimated β_3 than β_1 would support our prediction of stronger decrease in exploration over the business cycle in particular

---

[7] We aggregate to the 2-digit level to have enough observations per industry for a robust estimation.
[8] Because this is anyways not a true structural equation, it serves rather illustrative purposes, exactly as in Barlevy (2007). Coefficients are tightly estimated and not adjusted for estimation error.

for pro-cyclical industries. Again, we estimate the equation once with and without year fixed effects to allow an assessment of how industry specific cyclicality beyond the macroeconomic cycle as opposed to macroeconomic changes drive changes in innovative search. As a robustness check we further estimate the baseline model based on split samples, where we first focus on industries with a cyclicality measure below or equal to the median value as compared to particular pro-cyclical industries above the median value.

Table 4, columns (a) and (b), present the results of estimating (3), while columns (c) and (d) reflect the baseline results for particularly pro-cyclical industries only, and columns (e) and (f) reflect the corresponding other half of the sample. The results provide further support for our theoretical predictions. Firms tend to decrease their focus on exploration more sharply the stronger the cyclicality of the industry they operate in (an F-test of $\beta_1 - \beta_3 = 0$, is statistically significant at $p < 0.006$ (a) and $p < 0.04$ (b), respectively). In particular pro-cyclical industries we estimate that a one standard deviation increase in output corresponds to a -0.35 (model a, [-0.25, model b]) decrease in standard deviation of innovative search, while in weakly pro-cyclical and counter-cyclical industries, a one standard deviation increase in output corresponds to a -0.16 (model a, [-0.12, model b]) standard deviation decrease in innovative search.

Table 4 – Industry growth, innovative search and cyclicality

| | Innovative Search | | | | | |
| | Full sample | | Cyclicality > p50 | | Cyclicality <= p50 | |
	a	b	c	d	e	f
Log(R&D)$_{t-1}$	-0.000	-0.003	0.002	0.003	-0.004	-0.012**
	(0.003)	(0.003)	(0.003)	(0.004)	(0.006)	(0.006)
Log(Sales)$_{t-1}$	0.010*	0.012**	0.009	0.005	0.011	0.022**
	(0.005)	(0.006)	(0.006)	(0.007)	(0.010)	(0.009)
Log(Employees)$_{t-1}$	-0.040***	-0.056***	-0.048**	-0.061***	-0.029	-0.050**
	(0.014)	(0.014)	(0.020)	(0.020)	(0.020)	(0.020)
Log(Capital)$_{t-1}$	-0.027***	-0.018*	-0.024	-0.014	-0.030**	-0.024**
	(0.010)	(0.010)	(0.015)	(0.015)	(0.012)	(0.010)
Log(Output)$_{t-1}$	-0.033***	-0.025***	-0.069***	-0.045***	-0.032***	-0.024***
	(0.007)	(0.005)	(0.012)	(0.011)	(0.007)	(0.008)
Log(Output)$_{t-1}$ x Cyc	-0.066***	-0.051***				
	(0.012)	(0.012)				
N	21051	21051	8609	8609	12442	12442
Year fixed effects	No	Yes	No	Yes	No	Yes
Firm fixed effects	Yes	Yes	Yes	Yes	Yes	Yes
R^2	0.467	0.474	0.490	0.499	0.451	0.462

Notes: This table presents OLS regression of firms' innovative search focus, defined as the technological proximity between the patents filed in year t to the existing patent portfolio held by the same firm from year t-5 up to year t-1, calculated according to Jaffe (1989). Models c and d are only firms in industries where stock prices follow industry growth above median levels, while models e and f are only firms in industries where stock prices follow industry growth below or equal to the median level. The main effect of *Cyc* is fully absorbed by the firm fixed effects. Standard errors clustered at the industry level appear in parentheses. ***, ** and * indicate a significance level of 1%, 5%, and 10%, respectively.

4.3 Financial constraints

To test whether financially constrained firms are indeed less sensitive to downturns we split the sample according to firms' S&P credit ratings. The lower sample size results from the limited availability of credit ratings. Table 6, columns (a) and (b), present the results of estimating (2), where Cyc_k is replaced with a dummy indicating firms that had an investment grade rating on average over the sampling period. Columns (c) and (d) reflect the baseline results for firms with a speculative rating only, and columns (e) and (f) reflect the corresponding other half of the sample. Consistent with the prediction from theory, firms with an investment grade rating tend to decrease their focus on exploration more sharply over the business cycle. Financially constrained firms without a rating reduce their focus on exploration over the business cycle by -0.126 standard deviations (model a, [-0.095, model b]) per one standard deviation increase in industry output, while unconstrained firms reduce their focus about twice as much by -0.234

standard deviations (model a, [-0. 209, model b]) per one standard deviation increase in industry output.

If we proxy financial constraints by small firm size, we find consistent results (please see Appendix). Small firms tend to decrease their focus on exploration over the business cycle less sharply than large firms.

Table 5 – Financial constraints speculative vs investment grade firms

| | Innovative Search | | | | | |
| | Full sample | | Spec. grade firms | | Investment grade firms | |
	a	b	c	d	e	f
Log(R&D)$_{t-1}$	0.001	-0.002	-0.001	-0.009	0.002	0.003
	(0.003)	(0.004)	(0.006)	(0.007)	(0.004)	(0.005)
Log(Sales)$_{t-1}$	0.005	0.010	-0.010	-0.007	0.009	0.014
	(0.013)	(0.013)	(0.030)	(0.032)	(0.012)	(0.014)
Log(Employees)$_{t-1}$	-0.026	-0.049**	-0.039	-0.052	-0.014	-0.038
	(0.020)	(0.021)	(0.033)	(0.037)	(0.023)	(0.025)
Log(Capital)$_{t-1}$	-0.040***	-0.027**	-0.022	-0.018	-0.051***	-0.034**
	(0.012)	(0.012)	(0.019)	(0.023)	(0.016)	(0.015)
Log(Output)$_{t-1}$	-0.023**	-0.018**	-0.020*	-0.017	-0.044***	-0.039***
	(0.010)	(0.009)	(0.011)	(0.010)	(0.009)	(0.007)
Log(Output)$_{t-1}$ x Inv. Grade	-0.044***	-0.040***				
	(0.008)	(0.007)				
N	9568	9568	3491	3491	6077	6077
Year fixed effects	No	Yes	No	Yes	No	Yes
Firm fixed effects	Yes	Yes	Yes	Yes	Yes	Yes
R^2	0.486	0.498	0.468	0.493	0.476	0.487

Notes: This table presents OLS regression of firms' innovative search focus, defined as the technological proximity between the patents filed in year *t* to the existing patent portfolio held by the same firm from year *t-5* up to year *t-1*, calculated according to Jaffe (1989). The reported R^2 is the within firm explained variation. Standard errors clustered at the industry level appear in parentheses. ***, ** and * indicate a significance level of 1%, 5%, and 10%, respectively.

5. Discussion

While simple, the model remains consistent with the organizational realities of high technology firms. Such firms experience manufacturing and logistics pressures during booms as they respond to demand. Particularly in a crisis (for example, inordinate sales demand or a yield crash), managers of manufacturing will seek additional resources -- and the research and development organizations provide tempting repositories of highly talented and immediately effective help. Rather than increase head count and go through the laborious process of hiring and training new employees, a manufacturing manager will often prefer to request help from

his or her upstream functions. In a stable firm with low turnover, that manager will know and have worked with the same R&D engineers who invented and perfected the problematic product. Particularly during a yield or sales crisis, the R&D manager will find it difficult to avoid demands to help his or her manufacturing counterpart. Such temporary assignments will in turn delay exploration of new opportunities – and increase the firm's attention on current technologies.

Again consistent with the model, the pressures to siphon off exploration talent in order to fight immediate crises will be greater in cyclical industries, as for example, in semiconductors. Yield crashes in semiconductor fabs have myriad and interdependent causes, and often result from interactions between physical and process design (done in the R&D organization) and manufacturing implementation (done by the downstream organization). Unsolved problems can lead to cross functional accusations and the temporary re-assignment of R&D engineers to the fab floor, and that temporary re-assignment delays research.

The model's intuition behind financial constraint and decreased exploration can also be observed in how an executive and R&D manager chooses projects and product for development. At the extreme, when a firm sees the potential for bankruptcy, it will be an unusual manager who protects the long-term opportunities. Faced with severe pressure for revenue and immediate success in the product market, few managers will keep their more speculative projects funded.

Other realities are also consistent with the model and will drive the results reported here. Defensive patenting (Shapiro 2001) consolidates and protects market share and should rise when firms think that the cost and delay in patent pendency warrant the investment. This investment requires legal time and money and cannot ignore the non-trivial demand on inventors' time as well. Despite well-trained patent lawyers, inventors cannot avoid spending time in crafting and approving even incremental patents and this time distracts them from exploring new ideas and technologies. Firms also need to consider the delay in getting patent approval; patent "pendency" typically lasts one to three years. All of these costs are easier to justify with the expectation of a growing and robust market. In contrast, with a shrinking or stagnant market, searching for new markets becomes relatively more attractive.

This work joins a burgeoning literature that looks beyond R&D expenditure or patent and citation counts to measure different types and nuances of innovation (please see the Appendix for robustness checks using alternate measures of exploration and exploitation, including the number of new to the firm classes, backward, and self-citations). For example, Kelly et al. (2018) construct a quotient where the numerator compares a patent's lexical similarity to future patents and the denominator to past patents. This explicitly incorporates future development of successful search and novelty and clearly identifies technological pivots and breakthroughs. Patents which score highly on this metric correlate with future productivity of the economy, sector, and firm. Balsmeier, Fleming and Manso (2017) use several simple patent-based measures to show that independent boards shift a firm towards exploitation strategies. Akcigit and Kerr (2016) develop a growth model to analyze how different types of innovation contribute to economic growth and how the firm size distribution can have important consequences for the types of innovations realized.

The lag between research and patent application could in principle make it hard to find results with these more nuanced patent measures. If there is a long lead time from initiating research to patenting, we may not find countercyclical exploration in the patent application data even if firms were to start exploring new areas during recessions. However, Griliches (1990) finds that "patents tend to be taken out relatively early in the life of a research project," and that the lag between initial research and patent application is typically short.

This work investigated how macroeconomic conditions influence firms' innovation and in particular, the types of innovative search those conditions motivate within the firm. Future work could look at how search strategies influence profitability, growth, and productivity changes. For example, do exploitation strategies lead to short term profits and meager productivity improvement, and exploration to lagged profits and fundamental improvements? Can firms appropriate exploitation patents more easily, even though the gains are smaller? Conversely, are the gains larger with exploration patents, yet more likely to leak to competitors?

6. Conclusion

Schumpeter and others have argued that innovative activities should concentrate in recessions. However, using common measures of innovation, such as R&D expenditures and raw patent counts, previous research found that innovation is instead procyclical. We propose a solution to this puzzle by modelling innovative search as a within the firm tension between exploration and exploitation. We rely on changes in the distribution of a firm's patenting across new and old to the firm technology classes to separate exploration and exploitation. Consistent with the model, exploitation strategies are procyclical while exploration strategies are countercyclical. The results are stronger for firms in more cyclical industries and less financially constrained firms.

Taken together, these results point to potential costs related to pursuing macroeconomic stability as a policy goal. Macroeconomic fluctuations might be useful in promoting growth-enhancing exploration that would otherwise not take place in the economy. Further investigation on this issue could be fruitful.

References

Aghion, Philippe and Gilles Saint Paul, 1998. "Virtues of Bad Times: Interaction between Productivity Growth and Economic Fluctuations" Macroeconomic Dynamics, September, 2(3), p322-44.

Aghion, Philippe, Philippe Askenazy, Nicolas Berman, Gilbert Cette, 2012. "Credit Constraints and the Cyclicality of R&D Investment: Evidence from France" Journal of the European Economic Association 10(5), p1001-1024.

Akcigit, Ufuk and William Kerr, 2016. "Growth Through Heterogenenous Innovations", Journal of Political Economy (forthcoming).

Arrow, Kenneth, 1969, Classificatory notes on the production and diffusion of knowledge, American Economic Review 59, 29–35.

Balsmeier, Benjamin, Lee Fleming, and Gustavo Manso, 2017. "Independent Boards and Innovation," Journal of Financial Economics vol. 123, 536-557.

Barlevy, Gadi, 2007. "On the Cyclicality of Research and Development" American Economic Review, 97(4), p1131-1164.

Bloom, N., Schankerman, M., Van Reenen, J, 2013. "Indentifying Technology Spillovers and Product Market Rivalry" Econometrica, 81(4), 1347-1393.

Canton, Eric and Harald Uhlig, 1999. "Growth and the Cycle: Creative Destruction versus Entrenchment" Journal of Economics, 69(3), p239-66.

Comin, Diego and Mark Gertler, 2006. "Medium-Term Business Cycles" American Economic Review, September, 96(3), June, p523-51.

Cooper, Russell and John Haltiwanger, 1993. "The Aggregate Implications of Machine Replacement: Theory and Evidence" American Economic Review, June, 83(3), p181-186.

Chamberlain, G., 1986, "Asymptotic Efficiency in Semi-Parametric Models with Censoring," Journal of Econometrics, 32, 189–218.

Fatas, Antonio. 2000. "Do Business Cycles Cast Long Shadows? Short-Run Persistence and Economic Growth." Journal of Economic Growth, 5(2): 147–62.

Fabrizio, K. and U. Tsolmon 2014. "An empirical examination of the procyclicality of R&D investment and innovation." *The Review of Economics and Statistics* 96(4):662-675.

Field, A., 2003, "The Most Technologically Progressive Decade of the Century." American Economic Review, 93(4): 1399-1413.

Francois, P. and H. Lloyd-Ellis 2003. "Animal Spirits through Creative Destruction." American Economic Review 93(3): 530-50.

Geroski, Paul A., and C. F. Walters. 1995. "Innovative Activity over the Business Cycle." Economic Journal, 105(431): 916–28.

Griliches, Zvi. 1990. "Patent Statistics as Economic Indicators: A Survey." *Journal of Economic Literature*, 28(4): 1661–1707.

Hall, B. Griliches, Z., and Hausman, 1986. "Patents and R and D: Is There a Lag?" International Economic Review, vol. 27, issue 2, 265-83.

Kogan, L., Papanikolaou, A., Seru, A., Stoffman, N. 2017. Technological Innovation, Resource Allocation and Growth. Forthcoming: Quarterly Journal of Economics.

Kopytov, A. and N. Roussanov, M. Taschereau-Dumouchel, 2018. "Short-run pain, long-run gain? Recessions and technological transformation." NBER Working Paper 24373.

Manso, Gustavo. 2011. "Motivating Innovation." Journal of Finance, 66(5), p1823-1860.

March, James. 1991. "Exploration and Exploitation in Organizational Learning" Organization Science, 2(1), p71-87.

Pakes, A. and Z. Griliches, 1980. "Patents and R&D at the firm level: A first report." *Economics Letters*, vol. 5, issue 4, 377-381.

Rafferty, Matthew C. 2003. "Do Business Cycles Influence Long-Run Growth? The Effect of Aggregate Demand on Firm-Financed R&D Expenditures." Eastern Economic Journal, 29(4): 607–18.

Schumpeter, J. 1939. *Business Cycles: A Theoretical, Historical, and Statistical Analysis of the Capitalist Process.* New York, Mcgraw Hill.

Schleifer, A. 1986. "Implementation Cycles." Journal of Political Economy, 94(6): 1163-90.

Shapiro, C. 2001. "Navigating the Patent Thicket: Cross Licenses, Patent Pools, and Standard-Setting". In Jaffe, Adam B.; et al. Innovation Policy and the Economy. I. Cambridge: MIT Press. pp. 119–150.

Taleb, N. 2012. *Antifragile: Things That Gain from Disorder*. Random House Trade Paperbacks. New York.

Walde, Klaus, and Ulrich Woitek. 2004. "R&D Expenditure in G7 Countries and the Implications for Endogenous Fluctuations and Growth." Economics Letters, 82(1): 91–97

Wooldridge, J. M., 2002, Econometric Analysis of Cross Section and Panel Data, Cambridge, MA: The MIT Press.

Appendix

Here we present tables that report a wide variety of robustness checks, alternate measures, and deeper analyses:

- A1: Financial constraints measured as small vs large firms

- A2: Alternative measure of innovative search – fraction of new to the firm patents

- A3: Alternative measures of innovative search – backward and self-citations

- A4: Controlling for industry specific trends

- A5: Forward term and confoundedness test

- A6: Intensive vs extensive margin

- A7: Excluding first 5 years of firms' patenting activity

- A8: 2 year moving averages

- A9: Assuming exploration in periods of no patenting

- A10: Excluding firms with little patenting activity

- A11: Limiting analysis from 1958 to 1999

- A12: Aggregated industry measure

- A13: 3-digit-SIC aggregation

- A14: HHI control for competition

- A15: OLS without fixed effects

- A16: Graphical test for linearity

Table A1 splits the sample according to firm size as measured by total assets. Columns a and b present the results of estimating equation (2), where Cyc_k is replaced with an indicator for large firms (equal or above median size). Columns c and d reflect the baseline results for small firms only (below median size), and columns e and f reflect the corresponding other half of the sample. Large firms tend to decrease their focus on exploration more sharply over the business cycle; assuming such firms are less financially constrained, this is consistent with prediction and Table 5.

Table A1 – Financial constraints measured as small vs large firms

| | Innovative Search | | | | | |
| | Full sample | | Small firms | | Large firms | |
	a	b	c	d	e	f
Log(R&D)$_{t-1}$	-0.001	-0.004	-0.019**	-0.019***	0.001	-0.000
	(0.003)	(0.003)	(0.007)	(0.007)	(0.003)	(0.004)
Log(Sales)$_{t-1}$	0.010*	0.012**	0.009*	0.013**	0.009	0.007
	(0.005)	(0.006)	(0.005)	(0.005)	(0.017)	(0.017)
Log(Employees)$_{t-1}$	-0.036**	-0.053***	-0.015	-0.032	-0.025	-0.048**
	(0.014)	(0.014)	(0.029)	(0.032)	(0.019)	(0.019)
Log(Capital)$_{t-1}$	-0.025***	-0.017*	-0.015	-0.013	-0.038***	-0.020
	(0.009)	(0.009)	(0.010)	(0.010)	(0.014)	(0.014)
Log(Output)$_{t-1}$	-0.032***	-0.024***	-0.027**	-0.008	-0.033***	-0.028***
	(0.008)	(0.007)	(0.010)	(0.010)	(0.012)	(0.009)
Log(Output)$_{t-1}$ x Large	-0.036***	-0.027***				
	(0.008)	(0.006)				
N	21051	21051	10526	10526	10525	10525
Year fixed effects	No	Yes	No	Yes	No	Yes
Firm fixed effects	Yes	Yes	Yes	Yes	Yes	Yes
R^2	0.467	0.474	0.446	0.452	0.541	0.551

Notes: This table presents OLS regression of firms' innovative search focus, defined as the technological proximity between the patents filed in year *t* to the existing patent portfolio held by the same firm from year *t-5* up to year *t-1*, calculated according to Jaffe (1989). The main effect of *Large* is fully absorbed by the firm fixed effects. Standard errors clustered at the industry level appear in parentheses. ***, ** and * indicate a significance level of 1%, 5%, and 10%, respectively.

Our main explanatory variable, lagged aggregated industry output, should be exogenous to strategic choices of a single firm and not suffer from reverse causality. Endogeneity could still arise, however, from potential measurement error or unobserved time-varying firm characteristics that covary with industrial output and innovative search. In the absence of valid instruments, we cannot perfectly rule out concerns in this respect, but we try to address many potential threats by leveraging the panel structure and by running a variety of robustness and sensitivity checks. Table A2 provides an alternative measure of innovative search, with the fraction of patents in new to the firm tech classes. Firms appear less likely to patent in new to the firm classes during expansions.

Table A2 – Alternative measure of innovative search – fraction of new to the firm patents

	Fraction new patents	
	a	b
Log(R&D)$_{t-1}$	-0.565*	-0.016
	(0.317)	(0.354)
Log(Sales)$_{t-1}$	-1.189*	0.012
	(0.635)	(0.588)
Log(Employees)$_{t-1}$	-2.149	-3.018**
	(1.311)	(1.221)
Log(Capital)$_{t-1}$	-2.919***	-2.582***
	(0.815)	(0.911)
Log(Output)$_{t-1}$	-3.350***	-1.478***
	(0.702)	(0.458)
N	21051	21051
Year fixed effects	No	Yes
Firm fixed effects	Yes	Yes
R^2	0.357	0.364

Notes: This table presents OLS regression of firms' innovative search focus, defined as the fraction of patents filed in year *t* that are assigned to original USPTO tech class where the given firm has not patented previously. Standard errors clustered at the industry level appear in parentheses. ***, ** and * indicate a significance level of 1%, 5%, and 10%, respectively.

In addition, we re-estimated the baseline model using the amount of backward citations and self-backward citations, respectively, as the dependent variable, instead of the Jaffe measure. Increased backward citations indicate a more crowded space in prior art and self-citations indicate that a firm is building upon existing technologies, rather than exploring new areas. Both measures correlate with a broad battery of exploitation measures (Balsmeier, Fleming, and Manso 2017). Consistent with a decreased focus on exploration over the business cycle, Table 3 illustrates increased rates of backward and self-backward citations during expansions.

Table A3 – Alternative measures of innovative search – backward and self-citations

	Backward citations		Self-back-citations	
	a	b	c	d
Log(R&D)$_{t-1}$	0.295***	0.087***	0.319***	0.120***
	(0.028)	(0.024)	(0.034)	(0.031)
Log(Sales)$_{t-1}$	0.166**	0.035	0.269***	0.093*
	(0.074)	(0.033)	(0.092)	(0.052)
Log(Employees)$_{t-1}$	0.235	0.657***	0.174	0.541***
	(0.164)	(0.097)	(0.191)	(0.127)
Log(Capital)$_{t-1}$	0.365***	0.178***	0.418***	0.237***
	(0.063)	(0.049)	(0.068)	(0.060)
Log(Output)$_{t-1}$	0.579***	0.189***	0.663***	0.224***
	(0.148)	(0.033)	(0.156)	(0.043)
N	21051	21051	21051	21051
Year fixed effects	No	Yes	No	Yes
Firm fixed effects	Yes	Yes	Yes	Yes
R^2	0.680	0.741	0.695	0.733

Notes: This table presents OLS regression of firms' of the log of firms backward citations +1 (models a and b) and the log of firms back citations to own patents (models c and d). Standard errors clustered at the industry level appear in parentheses. ***, ** and * indicate a significance level of 1%, 5%, and 10%, respectively.***, ** and * indicate a significance level of 1%, 5%, and 10%, respectively.

Table A4 illustrates how results are robust to adding linear or log-linear industry specific trends, which should ameliorate concerns that the results are driven by secular trends.

Table A4 – Controlling for industry specific trends

	Innovative search			
	a	b	c	d
Log(R&D)$_{t-1}$	-0.001	-0.006	-0.003	-0.006*
	(0.003)	(0.003)	(0.003)	(0.004)
Log(Sales)$_{t-1}$	0.010*	0.009	0.008	0.009
	(0.005)	(0.006)	(0.005)	(0.006)
Log(Employees)$_{t-1}$	-0.034**	-0.046***	-0.033**	-0.048***
	(0.015)	(0.015)	(0.015)	(0.014)
Log(Capital)$_{t-1}$	-0.025**	-0.018*	-0.025**	-0.017*
	(0.010)	(0.010)	(0.010)	(0.010)
Log(Output)$_{t-1}$	-0.037***	-0.033***	-0.041***	-0.032***
	(0.007)	(0.007)	(0.007)	(0.006)
N	21051	21051	21051	21051
Year fixed effects	No	Yes	No	Yes
Firm fixed effects	Yes	Yes	Yes	Yes
R^2	0.473	0.480	0.474	0.481

Notes: This table presents OLS regression of firms' innovative search focus, defined as the technological proximity between the patents filed in year *t* to the existing patent portfolio held by the same firm from year *t*-5 up to year *t*-1, calculated according to Jaffe (1989). Models a and b estimated including 3-digit-SIC linear trends and models c and d are estimated including 3-digit-SIC log-linear trends. ***, ** and *indicate a significance level of 1%, 5%, and 10%, respectively.

Adding a forward term of industrial output to our model further shows that all explanatory power comes from the one year lagged industry output, suggesting that we are unlikely to pick up any unobserved trends (Table A5, columns a and b). We also ran models just with time fixed effects to eliminate potential confoundedness from endogenous control variables (Table A5, columns c and d).

Table A5 – Forward term and confoundedness test

	Innovative search			
	a	b	c	d
Log(R&D)$_{t-1}$	-0.002	-0.004		
	(0.003)	(0.003)		
Log(Sales)$_{t-1}$	0.004	0.010		
	(0.006)	(0.006)		
Log(Employees)$_{t-1}$	-0.040***	-0.056***		
	(0.015)	(0.015)		
Log(Capital)$_{t-1}$	-0.028***	-0.019**		
	(0.009)	(0.009)		
Log(Output)$_{t+1}$	-0.006	0.009		
	(0.010)	(0.010)		
Log(Output)$_{t-1}$	-0.031***	-0.032***	-0.057***	-0.035***
	(0.011)	(0.009)	(0.012)	(0.009)
N	19973	19973	21051	21051
Year fixed effects	No	Yes	No	Yes
Firm fixed effects	Yes	Yes	Yes	Yes
R^2	0.472	0.478	0.461	0.469

Notes: This table presents OLS regression of firms' innovative search focus, defined as the technological proximity between the patents filed in year t to the existing patent portfolio held by the same firm from year t-5 up to year t-1, calculated according to Jaffe (1989). ***, ** and * indicate a significance level of 1%, 5%, and 10%, respectively.

Table A6 splits the data into old vs. young firms (</> 26 years in the data), reveals no significant differences, and implies that the results are not driven by differences in sample composition over time.

Table A6 – Intensive vs extensive margin

	Firms < 26 year of data		Firms >= 26 years of data	
	a	b	c	d
Log(R&D)$_{t-1}$	-0.008	-0.009	0.002	-0.003
	(0.006)	(0.006)	(0.003)	(0.004)
Log(Sales)$_{t-1}$	0.009	0.011*	0.002	0.012
	(0.005)	(0.006)	(0.020)	(0.023)
Log(Employees)$_{t-1}$	-0.028	-0.033	-0.023	-0.047**
	(0.021)	(0.021)	(0.019)	(0.021)
Log(Capital)$_{t-1}$	-0.019	-0.012	-0.050***	-0.039**
	(0.011)	(0.011)	(0.016)	(0.016)
Log(Output)$_{t-1}$	-0.041***	-0.024***	-0.027***	-0.022***
	(0.010)	(0.007)	(0.008)	(0.008)
N	13285	13285	7766	7766
Year fixed effects	No	Yes	No	Yes
Firm fixed effects	Yes	Yes	Yes	Yes
R^2	0.457	0.465	0.473	0.485

Notes: This table presents OLS regression of firms' innovative search focus, defined as the technological proximity between the patents filed in year *t* to the existing patent portfolio held by the same firm from year *t*-5 up to year *t*-1, calculated according to Jaffe (1989). Standard errors clustered at the industry level appear in parentheses. ***, ** and * indicate a significance level of 1%, 5%, and 10%, respectively.

Results are further robust to excluding the first five years of patenting by each firms, where our measure might be particular noisy (see Table A7).

Table A7 – Excluding first 5 years of firms' patenting activity

	Innovative search	
	a	b
Log(R&D)$_{t-1}$	-0.000	-0.003
	(0.003)	(0.004)
Log(Sales)$_{t-1}$	0.017**	0.020**
	(0.008)	(0.008)
Log(Employees)$_{t-1}$	-0.040**	-0.060***
	(0.016)	(0.016)
Log(Capital)$_{t-1}$	-0.030***	-0.017*
	(0.009)	(0.009)
Log(Output)$_{t-1}$	-0.038***	-0.028***
	(0.011)	(0.008)
N	15728	15728
Year fixed effects	No	Yes
Firm fixed effects	Yes	Yes
R^2	0.486	0.494

Notes: This table presents OLS regression of firms' innovative search focus, defined as the technological proximity between the patents filed in year *t* to the existing patent portfolio held by the same firm from year *t*-5 up to year *t*-1, calculated according to Jaffe (1989). Standard errors clustered at the industry level appear in parentheses. ***, ** and * indicate a significance level of 1%, 5%, and 10%, respectively.

The same holds if we consider 2-year moving averages of our innovative search score, in order to reduce noise (see Table A8).

Table A8 – 2 year moving averages

	Innovative search	
	a	b
Log(R&D)$_{t-1}$	-0.001	-0.003
	(0.003)	(0.003)
Log(Sales)$_{t-1}$	0.010*	0.015**
	(0.006)	(0.006)
Log(Employees)$_{t-1}$	-0.042***	-0.060***
	(0.015)	(0.014)
Log(Capital)$_{t-1}$	-0.030***	-0.020**
	(0.009)	(0.009)
Log(Output)$_{t-1}$	-0.037***	-0.025***
	(0.008)	(0.006)
N	18790	18790
Year fixed effects	No	Yes
Firm fixed effects	Yes	Yes
R^2	0.606	0.615

Notes: This table presents OLS regression of firms' innovative search focus, defined as the 2-year moving average of technological proximity between the patents filed in year *t* to the existing patent portfolio held by the same firm from year *t*-5 up to year *t*-1, calculated according to Jaffe (1989). Standard errors clustered at the industry level appear in parentheses. ***, ** and * indicate a significance level of 1%, 5%, and 10%, respectively.

As we cannot observe firm search behavior when the firm does not patent, we assume they explore completely (innovative search is set to 1). This should alleviate concerns that our results are driven by the reduction of the sample to firm-year observations in which firms file at least one patent (see Table A9).

Table A9 – Assuming exploration in periods of no patenting

	Innovative search	
	a	b
Log(R&D)$_{t-1}$	-0.014***	-0.015***
	(0.003)	(0.003)
Log(Sales)$_{t-1}$	0.006**	0.001
	(0.003)	(0.003)
Log(Employees)$_{t-1}$	-0.053***	-0.060***
	(0.010)	(0.010)
Log(Capital)$_{t-1}$	-0.023***	-0.019***
	(0.004)	(0.004)
Log(Output)$_{t-1}$	-0.019***	-0.023***
	(0.005)	(0.005)
N	55681	55681
Year fixed effects	No	Yes
Firm fixed effects	Yes	Yes
R^2	0.640	0.646

Notes: This table presents OLS regression of firms' innovative search focus, defined as the technological proximity between the patents filed in year *t* to the existing patent portfolio held by the same firm from year *t*-5 up to year *t*-1, calculated according to Jaffe (1989). Standard errors clustered at the industry level appear in parentheses. ***, ** and * indicate a significance level of 1%, 5%, and 10%, respectively.

Filing only a few patents in a given year may create measurement issues, too, so we also estimate models where we restrict the sample to firm-year observations where firms applied for at least 2/5/10 patents, and found the results to be stable (see Table A10).

Table A10 – Excluding firms with little patenting activity

| | Innovative Search | | | | | |
| | Min 2 Patents | | Min 5 Patents | | Min 10 Patents | |
	a	b	c	d	e	f
Log(R&D)$_{t-1}$	-0.002	-0.004	-0.000	-0.001	0.000	-0.000
	(0.003)	(0.003)	(0.003)	(0.003)	(0.002)	(0.003)
Log(Sales)$_{t-1}$	0.003	0.007	0.007	0.013*	0.005	0.010
	(0.005)	(0.005)	(0.008)	(0.007)	(0.011)	(0.011)
Log(Employees)$_{t-1}$	-0.031**	-0.049***	-0.021*	-0.044***	-0.006	-0.029**
	(0.012)	(0.012)	(0.012)	(0.013)	(0.013)	(0.015)
Log(Capital)$_{t-1}$	-0.027***	-0.016*	-0.032***	-0.020**	-0.038***	-0.025**
	(0.008)	(0.008)	(0.009)	(0.009)	(0.010)	(0.010)
Log(Output)$_{t-1}$	-0.030***	-0.020***	-0.026***	-0.018**	-0.020***	-0.017***
	(0.008)	(0.006)	(0.009)	(0.008)	(0.007)	(0.006)
N	17738	17738	11932	11932	8163	8163
Year fixed effects	No	Yes	No	Yes	No	Yes
Firm fixed effects	Yes	Yes	Yes	Yes	Yes	Yes
R^2	0.497	0.506	0.535	0.545	0.558	0.568

Notes: This table presents OLS regression of firms' innovative search focus, defined as the technological proximity between the patents filed in year *t* to the existing patent portfolio held by the same firm from year *t*-5 up to year *t*-1, calculated according to Jaffe (1989). Standard errors clustered at the industry level appear in parentheses. ***, ** and * indicate a significance level of 1%, 5%, and 10%, respectively.

Results are also robust to taking out the 2000s years, which might be particularly influential due to the bust of the dotcom bubble (see Table A11).

Table A11 – Limiting analysis from 1958 to 1999

	Innovative search	
	a	b
Log(R&D)$_{t-1}$	0.001	-0.003
	(0.003)	(0.004)
Log(Sales)$_{t-1}$	0.004	0.013
	(0.008)	(0.008)
Log(Employees)$_{t-1}$	-0.022	-0.045***
	(0.014)	(0.015)
Log(Capital)$_{t-1}$	-0.041***	-0.034***
	(0.010)	(0.010)
Log(Output)$_{t-1}$	-0.040***	-0.028***
	(0.007)	(0.007)
N	15997	15997
Year fixed effects	No	Yes
Firm fixed effects	Yes	Yes
R^2	0.482	0.488

Notes: This table presents OLS regression of firms' innovative search focus, defined as the technological proximity between the patents filed in year *t* to the existing patent portfolio held by the same firm from year *t*-5 up to year *t*-1, calculated according to Jaffe (1989). Standard errors clustered at the industry level appear in parentheses. ***, ** and * indicate a significance level of 1%, 5%, and 10%, respectively.

Next, we exchanged the industry output measure with total shipments per sector as measured by the NBER productivity database (Table A12), and estimate models with higher aggregated output measures (3-digit-level, Table A13). The higher aggregation should lessen concerns that measurement error with respect to the relevant industries confound our results, because firms are active in more than one 4-dgit SIC industry.

Table A12 – Aggregated industry measure

	Innovative search	
	a	b
Log(R&D)$_{t-1}$	-0.002	-0.004
	(0.003)	(0.003)
Log(Sales)$_{t-1}$	0.009*	0.011*
	(0.005)	(0.006)
Log(Employees)$_{t-1}$	-0.040***	-0.056***
	(0.014)	(0.014)
Log(Capital)$_{t-1}$	-0.028***	-0.018*
	(0.010)	(0.009)
Log(Shipments)$_{t-1}$	-0.036***	-0.026***
	(0.008)	(0.006)
N	21051	21051
Year fixed effects	No	Yes
Firm fixed effects	Yes	Yes
R^2	0.466	0.474

Notes: OLS regression of firms' innovative search focus, defined as the technological proximity between the patents filed in year *t* to the existing patent portfolio held by the same firm from year *t*-5 up to year *t*-1, calculated according to Jaffe (1989). Standard errors clustered at the industry level appear in parentheses. ***, ** and * indicate a significance level of 1%, 5%, and 10%, respectively.

Table A13 – 3-digit-SIC aggregation

	Innovative search	
	a	b
Log(R&D)$_{t-1}$	-0.004*	-0.005
	(0.003)	(0.003)
Log(Sales)$_{t-1}$	0.006	0.010*
	(0.005)	(0.005)
Log(Employees)$_{t-1}$	-0.038***	-0.056***
	(0.013)	(0.013)
Log(Capital)$_{t-1}$	-0.027***	-0.017**
	(0.009)	(0.008)
Log(Output)$_{t-1}$	-0.029***	-0.015**
	(0.005)	(0.007)
N	29412	29412
Year fixed effects	No	Yes
Firm fixed effects	Yes	Yes
R^2	0.457	0.464

Notes: OLS regression of firms' innovative search focus, defined as the technological proximity between the patents filed in year *t* to the existing patent portfolio held by the same firm from year *t*-5 up to year *t*-1, calculated according to Jaffe (1989). Standard errors clustered at the industry level appear in parentheses. ***, ** and * indicate a significance level of 1%, 5%, and 10%, respectively.

Finally, some change in innovative search might be driven by changes in competition, so we ran additional models where we control for the sales based Herfindahl index per 4-digit SIC industry, once with a linear term and once with the squared term included (see Table A14). Our main results remain unchanged.

Table A14 – HHI control for competition

| | Innovative Search | | Innovative Search | |
	a	b	c	d
Log(R&D)$_{t-1}$	-0.002	-0.004	-0.002	-0.004
	(0.003)	(0.003)	(0.003)	(0.003)
Log(Sales)$_{t-1}$	0.009*	0.012**	0.009	0.012**
	(0.005)	(0.006)	(0.005)	(0.006)
Log(Employees)$_{t-1}$	-0.041***	-0.059***	-0.042***	-0.060***
	(0.014)	(0.014)	(0.014)	(0.014)
Log(Capital)$_{t-1}$	-0.028***	-0.018*	-0.027***	-0.018*
	(0.010)	(0.009)	(0.009)	(0.009)
HHI$_{t-1}$	0.041	0.055*	0.174*	0.171**
	(0.027)	(0.029)	(0.089)	(0.076)
HHI squared$_{t-1}$			-0.132	-0.116
			(0.085)	(0.075)
Log(Output)$_{t-1}$	-0.034***	-0.023***	-0.033***	-0.022***
	(0.008)	(0.006)	(0.009)	(0.007)
N	21051	21051	21051	21051
Year fixed effects	No	Yes	No	Yes
Firm fixed effects	Yes	Yes	Yes	Yes
R^2	0.467	0.474	0.467	0.475

Notes: This table presents OLS regression of firms' innovative search focus, defined as the technological proximity between the patents filed in year *t* to the existing patent portfolio held by the same firm from year *t*-5 up to year *t*-1, calculated according to Jaffe (1989). Standard errors clustered at the industry level appear in parentheses. ***, ** and * indicate a significance level of 1%, 5%, and 10%, respectively.

Results are also robust to models without firm fixed effects.

Table A15 – OLS without firm fixed effects

	Innovative search	
	a	b
Log(R&D)$_{t-1}$	-0.030***	-0.028***
	(0.004)	(0.004)
Log(Sales)$_{t-1}$	0.027***	0.026***
	(0.007)	(0.008)
Log(Employees)$_{t-1}$	-0.043***	-0.051***
	(0.010)	(0.010)
Log(Capital)$_{t-1}$	-0.021***	-0.019**
	(0.007)	(0.007)
Log(Output)$_{t-1}$	-0.040***	-0.019**
	(0.011)	(0.008)
N	21051	21051
Year fixed effects	No	Yes
Firm fixed effects	No	No
R^2	0.182	0.201

Notes: This table presents OLS regression of firms' innovative search focus, defined as the technological proximity between the patents filed in year *t* to the existing patent portfolio held by the same firm from year *t*-5 up to year *t*-1, calculated according to Jaffe (1989). Standard errors clustered at the industry level appear in parentheses. ***, ** and * indicate a significance level of 1%, 5%, and 10%, respectively.

Graphical test for linearity

Notes: This illustrates the relationship between log(industry output) and firms' innovative search focus, defined as the technological proximity between the patents filed in year *t* to the existing patent portfolio held by the same firm from year *t*-5 up to year *t*-1, calculated according to Jaffe (1989). The red line represents the relationship estimated with a standard OLS regression. For easier graphical inspection the data is sorted into 20 equal bins and each dot represent the mean of each bin.